

MULTIMOMENT RADIO TRANSIENT DETECTION

L. G. SPITLER, J. M. CORDES, S. CHATTERJEE

Astronomy Department and NAIC, Cornell University, Ithaca, NY, 14853

AND

J. STONE

Barnard College, New York, NY, 10027

Draft version January 17, 2012

ABSTRACT

We present a multimoment technique for signal classification and apply it to the detection of fast radio transients in incoherently dedispersed data. Specifically, we define a spectral modulation index in terms of the fractional variation in intensity across a spectrum. A signal whose intensity is distributed evenly across the entire band has a lower modulation index than a spectrum whose intensity is localized in a single channel. We are interested in broadband pulses and use the modulation index to excise narrowband radio frequency interference (RFI) by applying a modulation index threshold above which candidate events are removed. The technique is tested both with simulations and using data from known sources of radio pulses (RRAT J1928+15 and giant pulses from the Crab pulsar). The method is generalized to coherent dedispersion, image cubes, and astrophysical narrowband signals that are steady in time. We suggest that the modulation index, along with other statistics using higher-order moments, should be incorporated into signal detection pipelines to characterize and classify signals.

1. INTRODUCTION

Surveys have always played an important role in astronomy and will play an increasingly important role in the future as observatories such as the Large Synoptic Sky Telescope (LSST) and the Square Kilometer Array (SKA) come online. The huge volumes of data generated by surveys require robust pipelines to identify and characterize populations with maximal completeness and minimal false positives. Regardless of the target source class, all detection pipelines rely on signal-to-noise ratio (SNR) to find sources and quantify their believability. This approach underutilizes the spectral information contained in the data, because it only uses the total intensity, or first moment, of a spectrum. We propose adding a second statistic, the spectral modulation index, that uses both the first moment (signal mean) and second moment (signal variance) of a spectrum to classify signals found through their SNR.

Although the technique we present in this paper is applicable to any data collected as intensity versus time and frequency, we focus on surveys for fast radio bursts and use the modulation index to identify and remove narrowband radio frequency interference (RFI). Surveys at radio frequencies must contend with RFI because, if not mitigated or excised, it can produce many false positives in processing pipelines. One difficulty in removing RFI is that it arises from a wide variety of terrestrial sources with different signal characteristics. Moreover it can be episodic or simply transient in nature along with the astrophysical signals that we are interested in. RFI can be broad in time and narrow in frequency (e.g. Global Positioning System satellites) or broad in frequency and narrow in time (e.g. lightning). Some radar signals sweep in frequency and mimic the plasma dispersion of astrophysical bursts.

We define fast radio bursts as having characteristic widths less than about one second, so astrophysical plasma delays, such as those encountered in pulsar signals and in the class of transients known as rotating radio transients (RRATs, McLaughlin et al. 2006), are important. When applied to fast transients, our technique builds upon the methods presented in Cordes and McLaughlin (2003) and McLaughlin and Cordes (2003) and that led to the discovery of RRATs. However, the basic idea applies to transients of any duration.

Actual signals will blur the distinction between RFI and signals of interest because some RFI will be broadband and some broadband astrophysical signals will show significant frequency variation. Astrophysical sources may have a spectral dependence including a simple spectral index or stronger modulations like those seen in solar bursts. Compact sources of fast transients will typically show frequency modulation from interstellar scintillation. We consider these effects in our implementation of the method.

The modulation index is not a new statistic; it has been used to measure time variations in a variety of astronomical applications, such as variability of pulsars (Weisberg et al. 1986) and extragalactic sources (Kedziora-Chudczer et al. 2001) caused by interstellar scintillations. It has also been used to study properties of the solar wind (Spangler and Spitler 2004).

We lay out the mathematical groundwork for multimoment dedispersion and the calculation of the modulation index in § 2. In § 3 we discuss the implementation of the modulation index in a fast transient detection pipeline and present the results of a simulated single pulse detection pipeline. Also in § 3 we apply the technique to two known sources of single pulses: Crab giant pulses and RRAT J1928+15. The method is extended to other types of data sets in § 4. We discuss how the spectral

modulation index can be used to classify bursts from a variety of astrophysical sources in § 5 and make concluding remarks in § 6.

2. METHOD

For specificity we consider broadband astrophysical signals that are sampled as dynamic spectra; that is, a sequence of spectra separated in time by Δt_s with N_ν frequency channels spanning a total bandwidth B . Most of the cases we discuss in this paper will have time-bandwidth products well in excess of unity, i.e., $\Delta t_s B / N_\nu \gg 1$, as is consistent with fast-dump spectrometers used in surveys for pulsars and radio transients. Our discussion will also concentrate on incoherent (post-detection) dedispersion, although we briefly discuss applications where coherent dedispersion is used. To illustrate the basic method, we consider only a simple sum over frequency to yield an intensity time series; later we will consider interstellar dispersion delays in the sum and usage of the dispersion effect in discriminating astrophysical signals from RFI.

We define the modulation index m_I as the normalized standard deviation of the intensity I across frequency (ν),

$$m_I^2 = \frac{\overline{I^2} - \bar{I}^2}{\bar{I}^2}, \quad (1)$$

where the first and second moments, \bar{I} and $\overline{I^2}$, respectively, are given by

$$\overline{I^n} = N_\nu^{-1} \sum_\nu I^n(\nu). \quad (2)$$

The modulation index, m_I , characterizes the distribution of signal power in spectrum. If the power is distributed evenly across the band, the variance is small, and the spectrum has a low modulation index. A broadband pulse with a flat spectrum and infinite signal-to-noise ratio (SNR) has $m_I = 0$. In the opposite extreme where the power is localized to single spectral channel, the modulation index is $m_I \rightarrow \sqrt{N_\nu - 1}$ with increasing SNR. These simple examples illustrate how broadband astrophysical transients can be discriminated from RFI, which often consists of narrow spikes in frequency accompanied by time variability, by requiring that the modulation index be less than some ceiling, $m_{I,\max}$.

2.1. Multimoment Dedispersion

The modulation index improves upon current detection schemes by including more information about the signal through the calculation of both the first and second moments. The signal processing algorithms used in the detection of short-duration radio transients (i.e., dedispersion and smoothing) must therefore be expanded to higher order moments. Although we only use the first and second moments in this paper, one could consider statistics that use higher order moments (e.g., skewness and kurtosis), so we introduce general, n^{th} -order expressions.

Radio pulses traveling through the interstellar medium are subject to frequency-dependent dispersion. Standard pulsar processing techniques remove the effects of dispersion by shifting the intensity in each frequency channel

in time according to the ν^{-2} dispersion relation and averaging to increase the signal-to-noise ratio (SNR).

The standard incoherent or post-detection approach for calculating a dedispersed time series is

$$\bar{I}(t, DM) = \frac{1}{N_\nu} \sum_\nu I(t + t_{DM}(\nu), \nu) \quad (3)$$

where $I(t, \nu)$ is the time-frequency intensity data, N_ν is the number of frequency channels, and $t_{DM}(\nu)$ is the delay at frequency ν for dispersion measure DM. Throughout we will denote an average in frequency with a bar over the variable. Survey data must be dedispersed using a set of trial DMs because the DM of the target sources are not known *a priori*, except in special applications, such as searches of globular clusters with previously known pulsars. A list of candidate pulses is defined by applying a minimum SNR threshold (SNR_{\min}) to the set of dedispersed time series.

Equation 3 implicitly weights all frequency channels equally. In general though there will be effects, both astrophysical and instrumental, for which optimal detection (i.e., maximum SNR) requires unequal weighting. For example broadband astrophysical sources have spectra with slopes characterized by a spectral index α , and maximum SNR occurs when the frequency channels are given some weighting w_ν that reflects α . Similarly removing instrumental effects, like bandpass subtraction, may introduce channel-dependent root-mean-square noise. The generalized expression for a weighted, dedispersed, first moment time series is

$$\bar{I}(t, DM; \alpha) = \frac{\sum_\nu w_\nu (\nu/\nu_0)^\alpha I(t + t_{DM}(\nu), \nu)}{\sum_\nu w_\nu}. \quad (4)$$

A full discussion of frequency weights is presented in Section 2.2.

Optimization in detection can also be made by considering the duration of the pulse relative to the time resolution of the data. The effective time resolution can be decreased by smoothing the data, and the maximum SNR occurs when the effective time resolution of the smoothed data matches that of the pulse (Cordes and McLaughlin 2003). The simplest method of smoothing adds adjacent time samples until all of the signal's power is in a single sample. In general any smoothing technique can be defined by applying smoothing weights w_t to the time-frequency data and averaging in time;

$$I_s(t, \nu) = \frac{\sum_{t'} w_{tt'} I(t', \nu)}{\sum_{t'} w_{tt'}}, \quad (5)$$

where $I_s(t, \nu)$ is the smoothed time-frequency data and $I(t', \nu)$ is the original time-frequency data. A dedispersed, smoothed time series $\bar{I}_s(t, DM)$ is calculated according to Equation 3 substituting I_s for I .

The calculation of the modulation index requires the second moment of the intensity time series. We make one final generalization by expanding Equations 4 and 5 to higher order moments. The weighted, dedispersed n^{th} -

moment time series is

$$\overline{I^n}(t, DM; \alpha) = \frac{\sum_{\nu} w_{\nu} [(\nu/\nu_o)^{\alpha} I(t + t_{DM}(\nu), \nu)]^n}{\sum_{\nu} w_{\nu}}. \quad (6)$$

The smoothed, dedispersed n^{th} -moment time series is

$$\overline{I_s^n}(t, DM) = \frac{1}{N_{\nu}} \sum_{\nu} \left[\frac{\sum_{t'} w_{tt'} I(t' + t_{DM}(\nu), \nu)}{\sum_{t'} w_{tt'}} \right]^n. \quad (7)$$

Because the spectral modulation index is a measure of the frequency structure, and not the temporal structure, the data are smoothed first in time, thereby consolidating the signal into a single spectrum, before it is squared and averaged. The modulation index for smoothed data is then

$$m_I^2 = \frac{\overline{I_s^2} - \overline{I_s}^2}{\overline{I_s}^2}. \quad (8)$$

We have specified separate expressions for the dedispersed intensity with frequency weights and smoothing weights solely for clarity; an optimal detection scheme would employ both. In the remainder of the paper we will explore the role of smoothing on the modulation index but keep the assumptions that $w_{\nu} = 1$ and $\alpha = 0$.

An alternative formulation to calculating the smoothed, dedispersed time series is to square the data before smoothing:

$$\overline{I_r^n}(t, DM) = \frac{1}{N_{\nu}} \sum_{\nu} \frac{\sum_{t'} w_{tt'} I^n(t' + t_{DM}(\nu), \nu)}{\sum_{t'} w_{tt'}}. \quad (9)$$

This approach captures both the time and frequency variation of the data encompassed by the summations. Throughout this paper we will focus on the spectral modulation index but will discuss this “time resolved” modulation index in Section 2.6 when we discuss the signature of incorrect dedispersion.

2.2. Frequency weights

As discussed in the previous section, weighting the frequency channels non-uniformly when calculating a dedispersed time series can improve the SNR of the detection. In general an astrophysical signal will have a power law spectrum

$$P_{\nu t} = P_{\nu_o t} \left(\frac{\nu}{\nu_o} \right)^{-\alpha_o} \quad (10)$$

where $P_{\nu t}$ is the pulse flux density at time t and channel ν , ν_o is a reference frequency, and α_o is the source’s spectral index. Frequency weights that yield the best SNR act to flatten the spectrum

$$w_{\nu, \alpha} = \left(\frac{\nu}{\nu_o} \right)^{\alpha} \quad (11)$$

where optimal detection occurs for $\alpha = \alpha_o$. Uniform spectral index frequency weights implicitly assumes a flat

spectral index of $\alpha = 0$. Pulsars have typical spectral indices of $\alpha_o = 1.6$ with significant variation ($\alpha_{o, min} = 0$ and $\alpha_{o, max} = 3$) (Lorimer et al. 1995). Processing with an implied spectral index of $\alpha = 0$ is only ideal for the minority of pulsars with the lowest observed spectral indices. Instead, applying frequency weights corresponding to the mean pulsar spectral index would increase the SNR with little additional computational cost.

Surveys for new classes of sources where α is unknown or surveys searching for extremely weak examples of a known population may warrant a search over spectral index. Dedispersing data using a set of trial spectral indices would increase the required computation by a factor equal to the number of trial spectral indices. Such a search may only be practical for offline post-processing.

Frequency weights can also reflect frequency-dependent noise variations caused by instrumental effects. In general we can define a signal model

$$I_{\nu t} = b_{\nu} (T_{\nu t} + g_{\nu t} P_{\nu t}),$$

where b_{ν} is the bandpass shape, $T_{\nu t}$ is the system temperature, $g_{\nu t}$ is the gain (e.g., K Jy⁻¹) and $P_{\nu t}$ is the pulsar or transient flux density. The quantity we are interested in is $P_{\nu t}$, but the quantity we measure is $I_{\nu t}$. Isolating P_{ν} requires removing the three frequency-dependent instrumental effects, which may introduce frequency-dependent rms noise. For example flattening the bandpass (b_{ν}) will result in higher noise at the band edges. For systems operating at low frequencies (~ 100 MHz) and large total bandwidths, the system temperature may vary across the band due to the strong frequency dependence of the sky brightness temperature. Finally gain variations will be both time and frequency dependent due to the particularities of the instrument. Again one chooses w_{ν} such that SNR is maximized, so if the additive noise can be modeled as Gaussian white noise, weighting the channels by the inverse of their variance (i.e., $w_{\nu} \propto 1/\sigma_{\nu}^2$) results in the maximum SNR.

To estimate the importance of considering frequency-dependent noise variations, we adopted a simple model for system temperature $T_{\nu t} = 20\text{K} + 10\text{K}(\nu/\nu_o)^{-2.7}$ and a signal with a non-zero spectral index. The resulting SNR after correcting for the spectral index and averaging over frequency was compared to the standard case that assumes a flat spectral index. At higher observing frequencies (~ 1 GHz) the improvement is small ($\sim 0.1\%$) for small or large bandwidths because of the low sky brightness temperature. At lower observing frequencies (~ 100 MHz), there is an improvement in SNR by as much as $\sim 80\%$ for wide bandwidths. This is particularly relevant for observatories like the Low Frequency Array (LOFAR, de Vos et al. 2009) and Murchison Widefield Array (MWA, Lonsdale et al. 2009).

A non-zero spectral index will increase the variance of a spectrum and therefore also increase the modulation index. This is undesirable because it would mistakenly imply a lower filling factor and may result in a true signal being flagged as RFI. The magnitude of the increase depends on the observation frequency (ν_o), bandwidth (B), and spectral index. For $B/\nu_o \sim 0.1$ the modulation index increase is of order ~ 0.1 for $\alpha \sim 3$. In the extreme case of $B/\nu_o \sim 1$, the modulation increase is of order ~ 1 for $\alpha \sim 3$. Current instruments have $B/\nu_o \sim 0.1$, so the increase in modulation index due to a source’s spectral

index is negligible. The trend is for new instruments to have larger bandwidths, so eventually the modulation increase will be significant enough that correcting for the spectral index becomes necessary. Furthermore in surveys that employ a spectral index search, the modulation index aids the analysis, because the trial spectral index closest to the true spectral index has the lowest modulation index.

2.3. Modulation Index

The modulation index is a quantitative measure of the patchiness, or modulation, of a spectrum. It differentiates signals whose power is distributed evenly across the band from those whose power is isolated to a few frequency channels. Broadband and narrowband signals have small and large modulation indices respectively. The level of modulation is parametrized by a frequency filling factor $f_\nu = W_\nu/N_\nu$ where W_ν is the number of channels in the spectrum that contains signal. Broadband astrophysical signals have a high filling factor ($f_\nu = 1$). RFI can be both narrowband or broadband with a filling factor ranging from $f_\nu = 1/N_\nu$ to 1. As this paper focuses on detecting broadband signals, the ultimate goal is to define a modulation index cutoff, $m_{I,\max}$, that will enable us to flag signals with low filling factors and vastly reduce the number of candidates created by a signal detection pipeline.

The analysis below assumes that the data have been searched for candidate signals by requiring a candidate to have an SNR_t larger than a minimum SNR (SNR_{\min}), and the modulation index is only calculated for these candidate signals. Because we assume that our data has zero mean, either through construction in the case of simulated data or through bandpass subtraction in the case of real data, this assumption assures that $\bar{I} > 0$ and m_I is well-defined. Furthermore, our interpretation of m_I assumes $\text{SNR}_{\min} > 1$, which is reasonable assumption since such a low SNR_{\min} would result in a deluge of events.

To derive a simple, analytical expression for the dependence of m_I on f_ν , we consider a spectrum with N_ν channels that contains two components: noise and signal. The noise is assumed to have a mean of zero and variance σ_x^2 . The signal fills W_ν channels with intensity A_i in channel i . While each A_i may be different, it will be useful to define \bar{A} , the average signal intensity over W_ν channels. The average signal intensity over the entire band is $f_\nu \bar{A}$. As we are usually more interested in signal-to-noise ratios (SNR), we define a single-channel SNR, $\text{SNR}_{\nu t} = \bar{A}/\sigma_x$, and time series SNR, $\text{SNR}_t = \sqrt{N_\nu} f_\nu \text{SNR}_{\nu t}$. The latter expression assumes that the standard deviation of the noise in the times series is $\sigma_x/\sqrt{N_\nu}$.

The modulation index for this simple signal model is

$$m_I^2 = \frac{N_\nu}{\text{SNR}_t^2} + \frac{m_A^2}{f_\nu} + \frac{1 - f_\nu}{f_\nu} \quad (12)$$

where we have introduced a separate signal modulation index, $m_A = \sigma_A/\bar{A}$. This modulation index characterizes the inherent frequency structure of a signal, and for this initial discussion we consider only signals with negligible structure ($m_A \ll 1$). In Sections 2.4 and 2.5 we discuss the effects of interstellar scintillation and pulsar self-noise and will introduce a non-zero m_A .




Figure 1. Time series SNR vs. modulation index for a simulated dataset with $N_\nu = 256$ containing noise (open circles), dispersed pulses (filled circles), and two sets of one-channel-wide RFI spikes (open squares). The solid horizontal line represents the applied intensity threshold of $\text{SNR}_{\min} = 3$. The vertical dashed line represents $m_{I,T} = 5.3$ as calculated for $\text{SNR}_t = 3$ using Equation 14. The dotted curve shows $m_{I,bb}$ as a function of SNR_t (Equation 13).

To explore the behavior of the modulation index as a function of SNR_t and f_ν , we will look at the extrema of f_ν : $f_\nu = 1$ and $f_\nu = 1/N_\nu$. It is important to note that Equation 12, as well as the limiting expressions defined below, represents the ensemble average values of the modulation index. Statistically they are the average value expected for a given combinations of SNR_t and f_ν , but noise in the data will cause variation in the actual values. The expressions below are also idealized in so far as we have set $m_A = 0$.

We also simulated a time-frequency data set containing broadband and ultra-narrowband signals added to Gaussian noise to accompany the discussion. It was processed according to standard pulsar processing techniques, and the results are shown in Figure 1. The data set has $N_\nu = 256$ and includes 100 broadband pulses ($f_\nu = 1$) with $W_t = 1$ and $\text{SNR}_t = 10$, two sets of 100 ultra-narrowband spikes ($f_\nu = 1/N_\nu$) with $W_t = 1$ and $\text{SNR}_t = 5$ and $\text{SNR}_t = 10$ respectively, and 10^6 noise-only spectra with zero mean and $\sigma_x^2 = 1$. We ignored dispersion for simplicity, and the time series was calculated using Equation 3 with $t_{DM}(\nu) = 0$. An intensity threshold was applied to the resulting time series with $\text{SNR}_{\min} = 3$ (solid horizontal line), and the modulation index was calculated for samples above threshold. The vertical dashed line and dashed curve are explained below.

A broadband signal ($f_\nu = 1$) with no inherent structure does not increase a spectrum's variance, so the modulation index depends only on the number of frequency channels and the signal's SNR_t

$$m_{I,bb} = \frac{\sqrt{N_\nu}}{\text{SNR}_t}. \quad (13)$$

Throughout we will see that the number of frequency channels scales the modulation index but does not change the relative magnitude (i.e., signals with larger f_ν have

lower m_I). As the number of channels from a single instrument generally remains constant, it is unimportant to the interpretation of a single data set but is important when comparing data from different instruments and choosing an appropriate $m_{I,\max}$.

Equation 13 reveals a direct relationship between SNR_t and m_I . The modulation index of a broadband pulse is not arbitrary; rather, it must fall along a curve proportional to $1/\text{SNR}_t$. Furthermore for a signal with a given SNR_t , the broadband modulation index is the lowest m_I the signal may have on average, as any $f_\nu < 1$ increases the frequency modulation of a signal and correspondingly its modulation index. In Figure 1 the broadband pulses (closed circles) cluster in a stripe centered at $\text{SNR}_t = 10$ and $m_{I,\text{bb}} = 1.6$. While each pulse has an inherent $\text{SNR}_t = 10$, the underlying noise in the data spreads the SNR_t and m_I of individual realizations about the ensemble average value ($m_{I,\text{bb}}$) and along the curve given by Equation 13 (dotted curve).

In practice there is a special SNR_t : SNR_{\min} , the minimum SNR constraint applied to the time series in the “thresholding” operation. The corresponding modulation index is given by Equation 13

$$m_{I,T} = \frac{\sqrt{N_\nu}}{\text{SNR}_{\min}}. \quad (14)$$

This modulation index represents the average maximum $m_{I,\text{bb}}$ that may exist in a set of thresholded candidates. In real data there will be spectra with SNR above the SNR_{\min} due to noise alone. The modulation index of thresholded noise follows the same relation as broadband pulses but with $\text{SNR}_t \sim \text{SNR}_{\min}$. Most of the events due to noise will cluster near $m_{I,T}$ and SNR_{\min} , but the rare stronger noise events will pepper the broadband curve toward larger SNR_t and lower m_I . In Figure 1 thresholded noise is shown as open circles, and only 100 points are plotted to reduce clutter. The weakest noise clusters near $\text{SNR}_{\min} = 3$ and $m_{I,T} = 5.3$, and the stronger noise climbs the $m_I \propto 1/\text{SNR}_t$ curve (dotted curve).

Because $m_{I,T}$ is on average the largest value of modulation index for a broadband signal in a candidate list, it is an upper limit to the choice of $m_{I,\max}$. Choosing a $m_{I,\max} > m_{I,T}$ would only return events from thresholded noise or RFI. The dashed vertical line in Figure 1 shows $m_{I,T}$ for $\text{SNR}_{\min} = 3$. If a modulation index upper limit is applied at $m_{I,\max} = m_{I,T}$, events to the left of the line are kept (all of the broadband pulses and about 85% of the thresholded noise), while all events to the right are dropped (all the ultra-narrowband spikes and about 15% of noise events). A larger SNR_{\min} would raise the solid line and move the dotted line to the left, which reduces the false alarm rate due to noise but limits one to detecting stronger pulses that are presumably rarer.

For the ultra-narrowband case where $f_\nu = 1/N_\nu$ and $\text{SNR}_t \rightarrow \infty$, the dependence on SNR_t drops out, and the average modulation reduces to

$$m_{I,s} = \sqrt{N_\nu - 1}. \quad (15)$$

Note that the modulation index for spiky signals depends only on the number of channels. Equation 15 also gives the upper limit on m_I as any $f_\nu > 1/N_\nu$ reduces the modulation and decreases the modulation index. Spiky RFI is illustrated in Figure 1 as two stripes of open squares




Figure 2. Modulation index vs. filling factor calculated using Equation 12 with $N_\nu = 256$. From top to bottom (thin to thick), curves are shown for four values of SNR_t : 3, 5, 10, 100.

centered at $\text{SNR}_t = 10$ and 5 respectively and $m_{I,s} = 16$. Like the broadband signals, the noise spreads the points about the ensemble average value, but clearly $m_{I,s}$ does not depend on SNR_t for narrowband signals.

The modulation index for intermediate filling factors at constant SNR_t must transition smoothly from $m_{I,\text{bb}}$ to $m_{I,s}$ as f_ν goes from 1 to $1/N_\nu$. The exact manner of the transition is given by Equation 12. Figure 2 illustrates the analytical variation of modulation index with filling factor for four values of the time series SNR ranging from $\text{SNR}_t = 3$ (top) to $\text{SNR}_t = 100$ (bottom) and $N_\nu = 256$. At low filling factors all curves tend toward $\sqrt{N_\nu}$, and at high filling factors the stronger the signal, the lower the modulation index. Most of the drop in modulation index happens at $f_\nu < 0.1$, suggesting that this technique can easily classify signals with filling factors less than about 10% but is less sensitive for signals with moderate to high filling factors.

In the above analysis we have assumed that a signal is either narrowband or broadband. Reality is messier, and we might have overlapping signals that are both narrowband and broadband. For example a pulse might occur at the same time as persistent narrowband RFI. To explore this case we adopt a simple model looking at the modulation index of a broadband pulse with no intrinsic frequency structure contaminated by an RFI spike that is one channel wide (i.e. the “ultra-narrowband” case described above). Equation 16 estimates the modulation index for this two-component case:

$$m_{I,\text{bb}+s} = \frac{\sqrt{N_\nu(1 + \text{SNR}_{t,s}^2)}}{\text{SNR}_{t,\text{bb}} + \text{SNR}_{t,s}}, \quad (16)$$

where $\text{SNR}_{t,\text{bb}}$ and $\text{SNR}_{t,s}$ are the time series SNR of the broadband pulse and RFI spike respectively. Note when $\text{SNR}_{t,s} \rightarrow 0$, the above equation reduces to the expression for broadband pulses (Equation 13), and when $\text{SNR}_{t,\text{bb}} \rightarrow 0$, the above equation reduces to the expression for ultra-narrowband spikes (Equation 15, up to the “−1”, which is negligible for large N_ν). When $\text{SNR}_{t,\text{bb}} \sim \text{SNR}_{t,s}$, we see that $m_{I,\text{bb}+s} \approx 0.5m_{I,s}$. If we choose

the modulation index cutoff to be $m_{I,T}$, this implies a $\text{SNR}_{\min} \leq 2$ to allow $m_{I,bb+s}$ to be above threshold, which is an unreasonably low threshold for most surveys. Furthermore using Equation 16, we can estimate that in order for $m_{I,bb+s} \geq m_{I,T}$, $\text{SNR}_{t,bb}/\text{SNR}_{t,s} > \text{SNR}_{\min} - 1$. For example if $\text{SNR}_{\min} = 5$, the SNR_t of the pulse must be four times larger than the SNR_t of the narrowband RFI. Using only the SNR_t and modulation index, we could easily miss a pulse if it occurs concurrent with strong, narrowband RFI. This finding that the modulation index is more sensitive to narrowband signals than broadband signals is consistent with Figure 2 and suggests our method is not a substitute for RFI excision techniques that identify persistent, strong, narrowband RFI from raw data.

2.4. Interstellar scintillations

Small scale density irregularities in the ionized interstellar medium (ISM) scatter and refract radio waves. Diffractive interstellar scintillations (DISS) and refractive interstellar scintillations (RISS) are observational phenomena seen in compact radio sources due to these irregularities. Scintillations are characterized by intensity variations with typical time (Δt_{DISS}) and frequency scales ($\Delta \nu_{\text{DISS}}$), and for DISS in the strong scattering regime, a time series is modulated as a random variable $g_{\text{DISS}}(t)$ with an exponential amplitude distribution and fractional intensity variations on the order of unity. More generally, DISS varies with both time and frequency ($g_{\text{DISS}}(t, \nu)$) with the diffraction timescale scaling as $\Delta t_{\text{DISS}} \sim \nu^{1.2}$ and bandwidth scaling as $\Delta \nu_{\text{DISS}} \sim \nu^{4.4}$. Also note that $\Delta \nu_{\text{DISS}}$ and the pulse broadening time (τ_d) are Fourier transform pairs and given by $2\pi\Delta \nu_{\text{DISS}}\tau_d = C_1$ (Cordes and Rickett 1998) where $C_1 \approx 1$. For fast transients we can safely assume that the duration of the pulse is much less than the diffraction time scale (at least for $\text{DM} \lesssim$ a few hundreds at $\nu = 1$ GHz), whereas the scintillation bandwidth may be of the same order as the channel resolution or bandwidth of a spectrometer. Frequency structure caused by DISS will increase the variance and thereby increase the modulation index of a signal.

Understanding the frequency structure of a spectrum due to strong scattering is clarified by defining three limits based on the relative sizes of $\Delta \nu$, B , and $\Delta \nu_{\text{DISS}}$ where $\Delta \nu$ is the width of a single frequency channel in a spectrum (Cordes et al. 2004). When the scintillation bandwidth is smaller than the channel bandwidth ($\Delta \nu_{\text{DISS}} \ll \Delta \nu$), the instrument effectively averages over several “scintles” and the intensity modulations are quenched. In this case $m_A \ll 1$, and the modulation index follows the expressions in Section 2.3. Similarly in the other extreme where there is only one scintle across the band ($0.2B < \Delta \nu_{\text{DISS}}$), the intensity variation is nearly flat over the entire bandwidth and $m_A \ll 1$. For the intermediate case where there are several distinct scintles across the band ($\Delta \nu \lesssim \Delta \nu_{\text{DISS}} \lesssim 0.2B$), the amplitude of each scintle is exponentially distributed and $m_A \approx 1$. As instruments become increasingly wide band and $B/\nu \sim 1$, this situation will become ever more relevant. Examples of these three cases can be seen in Cordes et al. (2004) for giant pulses from the Crab pulsar.

Connecting this to the discussion in Section 2.3, the re-

sulting average modulation index of a spectrum containing Gaussian noise and a broadband signal with exponentially distributed amplitudes is given by Equation 12 with $f_\nu = 1$ and $m_A = 1$. For weak signals the leading term of Equation 12 dominates and the spectral modulation index increases only slightly over that for a perfectly flat signal (i.e. Equation 13). For example, the modulation index of a spectrum with $N_\nu = 256$ and $\text{SNR}_{\min} = 5$ increases from $m_I = 3.2$ to $m_I = 3.35$. For strong signals (i.e. $\text{SNR}_t \gg \sqrt{N_\nu}$) the first term in Equation 12 becomes negligible and $m_I \approx m_A \approx 1$ for all $\text{SNR}_t \gg \sqrt{N_\nu}$.

2.5. Self Noise in the Pulsar Signal

Broadband pulsar signals have been modeled as amplitude modulated noise (Rickett et al. 1975) and as modulated, polarized shot noise (Cordes 1976; Cordes et al. 2004). The noise in these models corresponds to the emission over a broad range of radio frequencies while the modulation accounts for pulse structure. In this context, the modulation index of the signal is nonzero but is still smaller than the modulation indices expected from RFI. In the limit where the noise has Gaussian statistics, the intensity modulation index of polarized noise is

$$m_I^2 = m_{\text{ISS}}^2 + (1 + m_{\text{ISS}}^2)(1 + d_p^2)/2, \quad (17)$$

where m_{ISS} is the modulation index of frequency structure from DISS and d_p is the degree of polarization. The largest modulation is for $m_{\text{ISS}} = d_p = 1$ when $m_I = \sqrt{3}$. The observed modulation will be reduced if frequency structure from DISS is much broader than the total bandwidth B or if multiple pulse structures are averaged over in a single frequency channel of the spectrometer.

2.6. Signature of Incorrect Dedispersion

A bright pulse in a survey dedispersed with a large number of trial dispersion measures will yield events at many incorrect DMs in addition to the correct one. The true DM will return the largest SNR_t and narrowest pulse width on average. The residual pulse smearing from neighboring, incorrect DMs yields a lower SNR_t and wider pulse. The larger the DM error, the smaller the SNR is on average until the SNR drops below the threshold. This SNR–DM signature is one of the tests that a signal is a true astrophysical pulse and not RFI. The modulation index of a pulse also increases as the DM error increases.

The two-dimensional filling factor for a dispersed pulse dedispersed with a DM error of δDM is

$$f_{\nu,t} = \left[1 + \frac{|\Delta t_{\delta\text{DM}}|}{\Delta t_p} \right]^{-1}, \quad (18)$$

where $|\Delta t_{\delta\text{DM}}|$ is the absolute value of the residual dispersion smearing and Δt_p is the intrinsic width of the pulse in seconds. For a pulse that is perfectly dedispersed, $\Delta t_{\delta\text{DM}} = 0$ and $f_{\nu,t} = 1$, while for a large DM error, $\Delta t_{\delta\text{DM}} \rightarrow \infty$ and $f_{\nu,t} \rightarrow 0$.

Ideal matched filtering of a dispersed pulse, either in the absence of dedispersion or from residual smearing from incorrect dedispersion, smooths over the duration of the pulse’s dispersion sweep, consolidating the signal into a single time bin. The signal has a lower SNR_t than the pulse’s intrinsic SNR_t because $f_{\nu,t} < 1$, but the

one-dimensional filling factor is still $f_\nu = 1$. The spectral modulation index increases slightly due to the drop in SNR_t (Equation 13). For example, a hypothetical, unresolved Crab giant pulse detected in the data set described in Section 3.4.2 with an intrinsic $\text{SNR}_t = 30$ has $m_I = 0.75$ when dedispersed using the true DM of the Crab pulsar ($\text{DM} = 56.71 \text{ pc cm}^{-3}$) and $m_I \approx 4.3$ when dedispersed with a DM error of $\delta\text{DM} = 2 \text{ pc cm}^{-3}$.

The spectral modulation index does reflect the signature of incorrect dedispersion but only indirectly. A more powerful statistic would depend on $f_{\nu,t}$ rather than f_ν . This is accomplished by calculating an alternative modulation index ($m_{I,r}$) using the time resolved intensity moments given by Equation 9. For the same hypothetical Crab giant pulse described above, the resolved modulation index is the same as the spectral modulation index ($m_I = 0.75$) when the pulse is correctly dedispersed, but it increases to $m_{I,r} = 24.7$ for a DM error of $\delta\text{DM} = 2 \text{ pc cm}^{-3}$.

Combining the information provided by the spectral modulation index and resolved modulation index helps to identify spurious events from incorrect dedispersion. By first applying a cutoff in spectral modulation index, events are classified as either broadband or narrowband in frequency. The resolved modulation index of the broadband signals distinguish those that are also broad in time from those that are narrow in time.

2.7. Modulation Index Cutoff

The role of the spectral modulation index in a source detection pipeline is analogous to the SNR threshold (SNR_{\min}). A signal is first classified as interesting or not based on its SNR. For candidate signals that are strong enough, the application of a spectral modulation index cutoff ($m_{I,\max}$) classifies the candidate as interesting or not based on how broadband it is. Just as SNR_{\min} is applied to data automatically, the modulation index cutoff can be applied without human supervision.

The choice of $m_{I,\max}$ is influenced by the characteristics of the instrumentation and pipeline, as well as the importance of astrophysical effects such as diffractive interstellar scintillations. The upper limit to $m_{I,\max}$ is given by modulation index at the SNR threshold ($m_{I,T}$). As explained in Section 2.3, $m_{I,T}$ is the largest modulation index on average for a broadband pulse with a perfectly flat spectrum in a survey with an applied SNR minimum SNR_{\min} (Equation 14). Choosing $m_{I,\max} > m_{I,T}$ only yields candidates from thresholded noise or narrowband RFI. One exception is astrophysical signals with frequency structure, such as that caused by DISS. But as we showed in Section 2.4, the increase in the modulation index for weak signals is small (\sim a few percent), and one could increase $m_{I,\max}$ to allow for weak, scintillating signals with minimal increase in false positives.

Choosing $m_{I,\max} < m_{I,T}$ will reduce the number of false positives caused by thresholded noise and incorrectly dedispersed pulses, because it effectively applies a larger SNR_{\min} (Equation 13). The cost is reduced sensitivity, as on average only pulses with SNR_t exceeding than the larger effective SNR_{\min} will be below $m_{I,\max}$. But for surveys where interstellar scintillation may be important, a hard lower limit to the modulation index cutoff is set by the nature of exponential statistics; namely,

$m_{I,\max} = 1$, as described in Section 2.4.

Real signals will not always cleanly divide into broadband or narrowband. For example RFI could be marginally broadband, a real pulse could have strong, narrowband structure, or both narrowband and broadband signals could occur simultaneously. But for fast transients we also have an additional parameter: dispersion. Looking at the peak dispersion measure of a candidate signal can break the degeneracy between RFI and an astrophysical signal with the same modulation index.

2.8. Correlation Bandwidth

While the modulation index provides information about the degree of modulation in a spectrum, it does not provide any information about the shape of the modulation. A signal with amplitude \bar{A} localized in N adjacent bins has the same variance as a signal with the same amplitude whose power is distributed in N isolated bins across the band. We therefore need a new measure that reflects the distribution of a signal in frequency.

We define the characteristic bandwidth (B_c) to be the typical width of the frequency structure of signal. A broadband signal has $B_c = B$, and a narrowband signal has $B_c \sim \Delta\nu$. We define the fractional correlation bandwidth as a measure of the typical correlation scale compared to the total bandwidth

$$\text{FCB} = \frac{B_c}{B}. \quad (19)$$

The B_c can be estimated by calculating the correlation length from the autocorrelation (ACF) of a spectrum, which we define to be the half width half maximum (HWHM) of the first lobe of the ACF.

Figure 3 illustrates the technique for two simulated spectra with the same total intensity and variance. The left spectrum contains a Gaussian-shaped signal with a FWHM of 32 frequency channels. The right spectrum was generated by randomly swapping the channels in the left spectrum in groups of four, so that both spectra have the same total intensity, variance, and modulation index ($m_I \approx 1.5$). The bottom panels show the ACFs for the two spectra. To have the zero lag of the ACF equal to m_I^2 , the ACF was scaled and offset by $\text{ACF}/(N_\nu \bar{I}^2) - 1$, where \bar{I} is the mean of the spectrum. For the Gaussian spectrum, $\text{FCB} = 0.18$ and is consistent with the autocorrelation of a Gaussian with a FWHM=32 channels. For the spiky spectrum, $\text{FCB} = 0.03$ is consistent with a signal with spikes approximately four channels wide. Both spectra could be interesting astrophysical signals. The left signal could clearly be caused by an astrophysical process. The right signal could be indicative of strong interstellar scintillations. In any case the fractional correlation bandwidth provides another parameter that can be used to automatically classify a signal. Note the values for the fractional correlation bandwidth were calculated automatically along with the ACF and other statistics, suggesting the FCB could be implemented in an unsupervised pipeline.

3. APPLICATION

This section applies the techniques described above to the detection of fast radio transients. First we discuss




Figure 3. Fractional correlation bandwidth (FCB) for two example spectra. The top panels show two simulated spectra with the same total intensity and variance and $N_\nu = 128$. The left spectrum contains a Gaussian with a FWHM frequency width of 32 channels, and the right spectrum is a spiky spectrum generated by randomly swapping groups of four channels in the left spectrum. The autocorrelation functions (ACF) of the sample spectra are shown in the bottom panels. The ACF has been normalized by the square of the mean of the spectrum and the number of channels all minus 1. The characteristic bandwidth is calculated as the half width at half maximum of the first lobe of the ACF, and the corresponding FCB is given in the upper right corner of the lower panels.

one approach of incorporating the calculation of the spectral modulation index into a transient detection pipeline. To illustrate the technique we simulate a transient detection pipeline and apply the detection scheme to real data containing known transients. Our simulations and applications to real data show that a modulation index cutoff efficiently flags RFI and significantly reduces the number of candidates. Note our implementation starts with a list of pulse candidates generated in the standard manner and not on the raw data directly. A discussion of how the modulation index could be used to flag raw data in real time is discussed in Section 4.4. Also note our technique is independent of the number of beams (e.g. Arecibo L-band Feed Array¹) or stations (e.g. Very Long Baseline Array, Thompson et al. 2011) used to collect the data.

3.1. Implementation in a Detection Pipeline

The pipeline makes two passes over the data. The first pass defines a list of candidate events, and the second pass calculates the second moment and modulation index of these events. This two-pass approach is adopted for practicality and flexibility and will be explained in detail below. Also recall that the modulation index requires the first and second *central* moments, so the data must be bandpass subtracted before \bar{I} and \bar{I}^2 are calculated.

The data are first dedispersed as described by Equation 3 to generate a first moment time series. Surveys use a range of trial dispersion measures, each generating its own time series. These dedispersed time series are smoothed in time by applying a template bank of matched filters with different properties to account for

pulses of different widths and profiles. A list of candidates is defined by applying a minimum SNR threshold (SNR_{\min}) to these smoothed, dedispersed time series. In this paper we focus on two specific implementations of matched filtering: boxcar smoothing and clustering (friends-of-friends).

Boxcar smoothing convolves the data with a boxcar function of length W_t samples. The corresponding smoothing weights for Equation 5 are $w_{tt'} = 1$ for $t' = 0, \dots, W_t$. In practice boxcar matched filtering is implemented by iteratively summing adjacent time samples so $W_t = 2^{n_{sm}}$ where n_{sm} is the number of smoothing iterations. The details of this technique are described in Cordes and McLaughlin (2003). A single signal will likely be detected at several values of n_{sm} , so the event list should be sifted for the boxcar width that yields the maximum SNR.

The cluster, or friends-of-friends, algorithm looks for groupings of events in a time series. A cluster is defined as a set of events for which there is no gap in samples larger than N_{gap} . For a cluster with $N_{cluster}$ samples, the smoothing weights are $w_{tt'} = 1$ for $t' = t_1, \dots, t_N$ where t_1 is the first sample in a cluster, t_N is the last sample, and $t'_{i+1} - t'_i < N_{gap} + 1$. Importantly, the cluster algorithm is agnostic to the symmetry of the pulse, unlike the boxcar smoothing filter which is symmetric. As pulses from highly scattered sources have exponential tails, this algorithm may be better suited to detecting such astrophysical objects.

After a list of possible candidates is defined, one goes back to the raw time-frequency data to calculate the second moments and modulation indices. We grab a narrow range of raw data centered at the location of the candidate and reprocess it. The time-frequency snapshot is dedispersed and smoothed at the DM and smoothing parameters determined from the first pass. A time series is calculated for both the first and second moments. The first moment time series is thresholded in intensity, and the modulation index is calculated for the samples above threshold. The modulation indices of the events are compared to the modulation index cutoff, $m_{I,\max}$, and flagged as either a signal of interest if $m_I \leq m_{I,\max}$ or RFI if $m_I > m_{I,\max}$.

The two-pass approach is adopted out of practical considerations. While the first moment of the dedispersed time series can be smoothed directly, the time-frequency data must be smoothed before being squared and averaged in frequency (Equation 7). This would require a different dedispersed, smoothed, second moment time series for each matched filter type and parameter. Furthermore some smoothing approaches, such as the cluster algorithm, determine the smoothing weights $w_{tt'}$ from the thresholded first moment time series, making a parallel calculation of the second moment time series impractical.

3.2. Processing Requirements

For most cases the two pass approach is more computationally efficient than the obvious alternative of calculating the second moment in parallel with the first. We parameterize the processing required by the second pass in terms of the processing required to do the dedispersion in the first pass. Generally dedispersion dominates the processing time in a transient survey, so this is a useful metric. The number of operations required to

¹ <http://www.naic.edu/alfa>

dedisperse a block of time-frequency data with N_t time samples and N_ν frequency channels with N_{DM} trial dispersion measures is $N_{\text{ops},1} = N_t \times N_\nu \times N_{\text{DM}}$. If our first pass generates N_{events} candidate events, the number of operations required to dedisperse a narrow time range with N_s samples around each event at a single DM is $N_{\text{ops},2} = 4N_{\text{events}} \times N_s \times N_\nu$. The factor of 4 was included to consider the squaring of the data, summing of both the original and squared data and bandpass subtraction. The processing required by the second pass normalized by the dedispersion processing of the first pass is

$$P_2 = \frac{N_{\text{ops},2}}{N_{\text{ops},1}} = \frac{4N_{\text{events}}N_s}{N_{\text{DM}}N_t}. \quad (20)$$

If $N_t = 10^6$, $N_{\text{DM}} = 10^3$, $N_s = 10^3$, $N_{\text{events}} = 10^4$, P_2 is a few percent of the original dedispersion processing. This analysis ignores the additional overhead incurred by returning to the raw data, and in particular one must be wary of excess file I/O.

Most RFI excision techniques operate on raw data, not on the list of candidate pulses. Our technique is more general, because it can be used at both the beginning (see Section 4.4) and end of a source detection pipeline. The most common RFI excision approach applied to a list of candidates is to remove all events at low dispersion measure under the assumption that terrestrial signals are not dispersed. This is a blunt instrument and does not remove events from RFI at higher DMs. Calculating the modulation index of the candidate events allows for more sophisticated RFI excision in a list of candidates.

Classifying signals with the modulation index should be used in conjunction with other RFI excision algorithms. As described in Section 2.3, the modulation index is more sensitive to narrowband signals, and weak pulses may be missed if they occur simultaneously with strong, narrowband RFI. This suggests the modulation index algorithm works best together with algorithms that remove channels that contain persistent RFI. Similarly broadband RFI has a low modulation index, so removing impulsive RFI through other means will reduce the number of events from RFI that fall below the modulation index cutoff.

3.3. Simulations

To assess the usefulness of the modulation index as a signal diagnostic, we simulated a single-pulse event detection pipeline using Python-based software. In our simulations fake time-frequency data can be generated with Gaussian-distributed noise, dispersed pulses, and a variety of RFI. A dispersed pulse is added with a specified dispersion measure and Gaussian pulse profile with a FWHM of W_t . Spiky RFI is modeled as a two-dimensional Gaussian with a FWHM width in time (W_t) and frequency (W_ν). Broadband RFI is modeled as an undispersed pulse (i.e., $\text{DM}=0 \text{ pc cm}^{-3}$). The simulated data then undergo single pulse search processing. First the time-frequency data are dedispersed over a range of trial dispersion measures and both \bar{I} and \bar{I}^2 is calculated from the dedispersed time series. \bar{I} is thresholded, and for samples that are above threshold, the modulation index is calculated. Note that because these simulations involve a small amount of data, the two-pass analysis as described in Section 3.1 is not necessary, because all in-

termediate data products (i.e., dedispersed time series) can be kept in memory.

The fake data shown in panel (a) of Figure 4 set contains a single dispersed pulse, a Gaussian RFI spike, and a broadband RFI spike. The data properties are $N_\nu = 256$, $\Delta t_s = 1 \text{ ms}$, $N_t = 2000$ time samples, $\nu_o = 1400 \text{ MHz}$, and $B = 100 \text{ MHz}$. A single dispersed pulse was added at $t = 0.25 \text{ s}$ with $\text{DM} = 500 \text{ pc cm}^{-3}$, $\text{SNR}_{\nu t} = 1$, and $W_t = 1$. A narrowband Gaussian spike was added at $t = 1.0 \text{ s}$ and $\nu_o \approx 1428 \text{ MHz}$ with $\text{SNR}_{\nu t} = 40$, $W_t = 2$, and $W_\nu = 2$. Finally a broadband RFI spike is represented by an undispersed pulse (i.e., $\text{DM} = 0 \text{ pc cm}^{-3}$) at $t = 1.5 \text{ s}$ with $\text{SNR}_{\nu t} = 5$ and $W_t = 1$. The data were dedispersed over a range of trial dispersion measures $\text{DM} = 0\text{--}1000 \text{ pc cm}^{-3}$ and DM interval $\Delta\text{DM} = 6 \text{ pc cm}^{-3}$. A list of candidate event was defined by applying an SNR threshold of $\text{SNR}_{\text{min}} = 3$ to the resulting time series.

Figure 4 shows the results of this simulation. The top frame shows (a) the time-frequency data, (b) candidate events vs. DM and time, (c) $50/m_I$ for the candidate events, and (d) intensity vs. DM and time for candidates below $m_{I,\text{max}} = 3.2$. For clarity the SNR and pulse widths of the three signals in panel (a) have been exaggerated, and the noise is not shown. Instead of plotting the intensity, panel (b) shows a dot for each sample that was above the intensity threshold to reduce clutter. Panel (c) shows that most of the events have similar modulation indices; the one exception is for the broadband pulse, which has a modulation index almost an order of magnitude lower than all other points due to its high SNR. Panel (d) plots the SNR of the events with modulation indices below $m_{I,\text{max}}$ with the area of the circle proportional to the SNR of the event.

The triangle-shaped group of events near $t \sim 1.5 \text{ s}$ in panels (b) and (c) are spurious hits caused by the dedispersion path crossing the bright, broadband RFI samples. Most of these events have low filling factors because the RFI samples contribute only a few samples to the dedispersed spectrum and are not present in panel (d). The exception is low DM and $t = 1.5 \text{ s}$ where signal from the RFI contributes to many frequency channels resulting in a low modulation index. This is the incorrect dedispersion effect described in Section 2.6. While this example of RFI does pass our modulation index filter, the low DM of the event exposes it as RFI. Applying yet another filter that removes events at low DM will remove such events. The narrowband Gaussian spike causes a stripe of spurious events between $t \approx 0.9\text{--}1.0 \text{ s}$ as the path of the each trial dispersion measure crosses the spike. But for reasons just described above, these spurious candidates have modulation indices above our threshold. Finally the true astrophysical dispersed pulse, which is buried in the middle panels of Figure 4, stands out in panel (d) at $t = 0.25 \text{ s}$ after applying the modulation index cutoff.

The bottom frame in Figure 4 plots SNR_t vs. m_I and a histogram of m_I to illustrate how the modulation index groups signal types. The spurious events caused by the Gaussian spike (medium gray) are clumped together on the far right of the plot between $8 < m_I < 20$, consistent with the extreme narrowband case given by Equation 15 ($m_{I,s} = 16$). The events due to the dispersed

pulse and broadband RFI follow the light gray and black tracks respectively. The tracks are caused by incorrectly dedispersing the signal, and in both cases the lowest modulation index corresponds to the correct dispersion measure. The events from thresholded noise (black) are clumped near $m_{I,T} \sim 5$ and $\text{SNR}_t \sim \text{SNR}_{\min}$ as predicted by Equation 14. Because the weaker events associated with the broadband RFI overlap with the area containing thresholded noise, we've plotted them with the same color. The histogram of m_I in the lower panel shows that by applying a $m_{I,\max} = 3.2$, we have eliminated most of the events.

3.4. Application to Data

In the next two subsections we apply our modulation index classification technique to two known sources of single pulses detectable by single-pulse search pipelines. In Section 3.4.1 we apply the method to RRAT J1928+15 and in Section 3.4.2 to giant pulses from the Crab pulsar. Although in both cases we know the correct dispersion measure of the source, we re-analyze the data at a range of trial dispersion measures to recreate typical survey results.

The data were bandpass corrected by dividing by the median bandpass and subtracting off the mean. Dividing by the median spectrum flattens the spectra, assuring the bandpass shape is not contributing to the variance calculation. We choose the median because it is less sensitive to extreme values caused by pulses or RFI. Subtracting the mean assures us that our spectra have zero mean.

3.4.1. RRATs

RRAT J1928+15 was discovered by Deneva et al. (2009) in the PALFA survey running at the Arecibo Observatory. (See Deneva et al. 2009 for the details of the observation and PALFA parameters.) In brief, J1928+15 was discovered using the single-pulse search algorithms implemented by the Cornell pulsar search pipeline. Three pulses were detected with an interval of 0.403 s at $\text{DM} = 242 \text{ pc cm}^{-3}$. We reprocessed the bandpass corrected data with 642 trial dispersion measures roughly equally spaced over a DM range $\text{DM} = 0 - 500 \text{ pc cm}^{-3}$. The dedispersed time series had an intensity threshold applied at $\text{SNR}_{\min} = 4$. The samples above the SNR threshold were run through the cluster algorithm, and the sample in each cluster with the maximum intensity was used to calculate the modulation index.

The results are shown in Figure 5. In all panels only events with $\text{SNR} > 5$ are shown. The top frame plots intensity versus DM and time for all events above threshold (top) and the events that also satisfy $m_I < 1$. (Compare to Figure 3 in Deneva et al. 2009). The strong RRAT pulse is clearly visible around $t = 100 \text{ s}$, and a weak broadband RFI spike occurs near $t = 89 \text{ s}$. For the strong RRAT pulse we measure $\text{SNR}_t \approx 20$, and for the preceding weaker pulse, $\text{SNR}_t \approx 6$. The difference between our values and those in Deneva et al. (2009) is likely due to the fact that we used the cluster algorithm and they used the boxcar smoothing matched filter algorithm. We also do not detect the weaker, tailing pulse, likely for the same reason. Applying a $m_{I,\max} = 1$ eliminates $\approx 99\%$ of the events, leaving only the strongest events associated with the bright RRAT pulse.




Figure 4. Diagnostic plots for the simulated data described in Section 3.3. Top Frame: Panel (a) shows a grey scale of the simulated time-frequency data where the SNRs and widths of the signals have been exaggerated for clarity. Panel (b) plots a point in the DM-time plane for each event above the intensity threshold ($\text{SNR}_{\min} = 3$). Panel (c) plots $50/m_I$ for the candidate events in panel (b). Panel (d) plots the SNR of candidates over the SNR threshold and below the modulation index threshold ($m_{I,\max} = 3.2$). The data were dedispersed over a DM range of 0 to 1000 pc cm^{-3} with $\Delta\text{DM} = 6 \text{ pc cm}^{-3}$. The dispersed pulse is located at $\text{DM} = 500 \text{ pc cm}^{-3}$ and at $t = 0.5 \text{ s}$. A Gaussian RFI spike is located at $t = 1.0 \text{ s}$ and a frequency of $\sim 1428 \text{ MHz}$ and causes the stripe of events between $t \approx 0.9 - 1.0 \text{ s}$. Broadband RFI is located at $t = 1.5 \text{ s}$ and causes the triangle-shaped patch of candidate events. Bottom Frame: The top panel plots SNR_t vs. m_I , and the bottom panel is a histogram of modulation index. Events associated with the broadband RFI (and thresholded noise) are plotted in black, events associated with the Gaussian RFI spike are shown in medium gray, and the events associated with the dispersed pulse are plotted in light gray.

The bottom frame shows SNR_t vs. m_I (top) and a histogram of m_I (bottom). In the top panel the events associated with the strong RRAT pulse are localized along the upper curve extending from $\text{SNR}_t \approx 20$ and $m_I \approx 0.8$ to $\text{SNR}_t \approx 5$ and $m_I \approx 3$. The lower limit of $m_I \approx 0.8$ is consistent with the value predicted by Equation 13 for $\text{SNR}_t = 20$. The curve itself comprises spurious events caused by the strong pulse being dedispersed at incorrect dispersion measures and account for

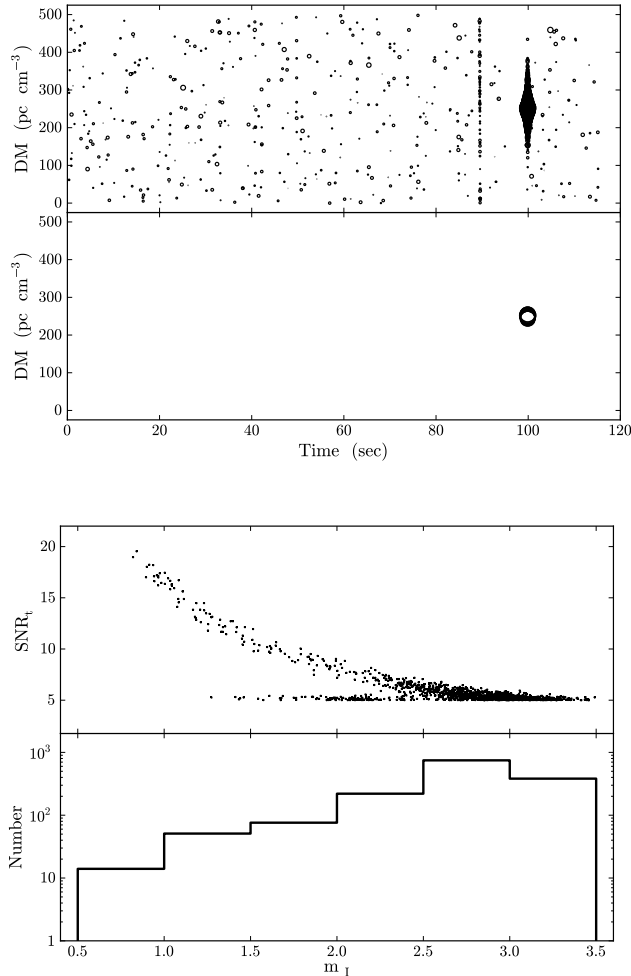


Figure 5. Modulation index calculation for the RRAT J1928+15. In all frames events with $\text{SNR} > 5$ are shown. Top Frame: Intensity vs. DM and time (top) for the events above the SNR threshold and the same events (bottom) with an additional modulation index threshold ($m_{I,\text{max}} < 1$) applied (bottom). The area of each circle is proportional to the SNR of the event. (See Deneva et al. 2009, Figure 3.) The RRAT pulse is clearly visible around $t = 100$ s and $\text{DM} = 242 \text{ pc cm}^{-3}$ and a weak, broadband RFI spike around $t = 89$ s. Bottom Frame: The SNR_t vs. m_I (top) and a histogram of m_I (bottom). The data were dedispersed using ~ 640 trial dispersion measures ranging from 0 to 500 pc cm^{-3} .

about 90% of all events in this figure. As described in Section 2.6, applying a modulation index cutoff can reduce the number of spurious events from real signals at incorrect DMs. The broadband RFI spike produced the events along the vertical line along $\text{SNR}_t = 5$ and $1.25 < m_I < 2.5$. The events associated with thresholded noise are grouped between $2.5 < m_I < 3.5$ and are consistent with $m_{I,T} = 3.2$ for $N_\nu = 256$ and $\text{SNR}_{\text{min}} = 5$. While applying $m_{I,\text{max}} = 1$ clearly isolated the bright pulse from the RRAT, we did so at the cost of applying an effective $\text{SNR}_{\text{min}} = 10$ and thereby removed the weaker RRAT pulse.

3.4.2. Crab giant pulses

The most well-known source of single, dispersed pulses is the Crab pulsar. We reprocessed a 430 MHz data set

containing the “supergiant” pulse with a $\text{SNR}_t \sim 1500$ reported in Cordes et al. (2004). The 140-second long dataset was dedispersed using 616 trial dispersion measures ranging from $\text{DM} = 0 - 250 \text{ pc cm}^{-3}$. An intensity threshold of $\text{SNR}_{\text{min}} = 4$ was applied to the dedispersed time series, and the samples above threshold were run through the cluster algorithm. The modulation index was calculated for the sample from each cluster that had the maximum intensity. This data set proved to be a particularly interesting case study due to the large number of pulses with a range of SNRs.

The results of the reprocessing are shown in Figure 6, and in all panels only events with $\text{SNR} > 5$ are shown. The top frame shows the intensity for the events versus DM and time with a circle whose area is proportional to the SNR of each event. The top panel shows all events, and the bottom panel shows those events that also fall below our modulation index cutoff ($m_{I,\text{max}} < 2$). A train of normal giant pulses is clearly visible along the dispersion measure of the Crab, $\text{DM} = 56.71 \text{ pc cm}^{-3}$, and the supergiant pulse is located near $t = 50$ s. This pulse is so bright that it has some of the same characteristics as broadband RFI; namely, events occurring at high dispersion measures. But as we saw previously for J1928+15, applying a $m_{I,\text{max}}$ eliminates many of the spurious events at incorrect DMs because they have low filling factors.

The bottom frame plots SNR_t vs. m_I (top) and a histogram of m_I (bottom). The events associated with the supergiant pulse are shown in medium gray and fall along the curve extending from $10^4 > \text{SNR}_t > 100$ and $0.7 < m_I < 11$ as well as the clump at $\text{SNR}_t \sim 10$ and $m_I \sim 15$. Note that the largest modulation indices are consistent with the value for spiky RFI ($m_{I,s} \approx 21$). The correctly dedispersed normal giant pulses are shown in black and lie along the left edge of the main cluster ranging from $100 > \text{SNR}_t > 5$ and $0.5 < m_I < 3$. One exception is the single, larger black point at the low- m_I end of the supergiant pulse curve (medium gray) showing the point associated with correct dedispersion value. Not surprisingly this event has the largest SNR_t and lowest m_I . The events from incorrectly dedispersing the regular giant pulses (light gray) spread to lower SNR_t and higher m_I than their correctly dedispersed counterparts (black). Finally the noise (light gray) falls between $2 \lesssim m_I \lesssim 7$ and along $\text{SNR}_t \sim 5$ and is consistent with $m_{I,T} \approx 4.2$.

The measured values of m_I for the Crab pulses are systematically larger than that predicted by Equation 13 for broadband pulses. The brightest regular giant pulses have $\text{SNR}_t \approx 100$ and $m_I \approx 0.5$, but the modulation index of a perfectly flat pulse with this SNR_t is ~ 0.2 . This suggests that Crab giant pulses have significant inherent frequency structure and non-zero m_A . Most striking is that the modulation index for the supergiant pulse ($m_I \approx 0.7$) is larger than the modulation index for the strongest regular pulses even though its SNR is two orders of magnitude larger. The spectrum of the supergiant pulse shows significant variation across the band due to DISS, as shown in Cordes et al. (2004).

The histogram of m_I shows the total number of events at all DMs (thick line) and the number of events at the DM of the Crab (thin line) as an proxy for the number of giant pulses. The assumption that any event at the DM of the Crab is a giant pulse is simplistic and may include false positives. Because brighter giant pulses are




Figure 6. Modulation index calculation for the Crab pulsar at 430 MHz. In all frames only events with $\text{SNR} > 5$ are shown. Top Frame: Events above the intensity threshold are plotted against DM and time (top) and with the additional constraint $m_{I,\text{max}} < 2$ (bottom). A stream of Crab giant pulses can be seen in the DM-time plot at $\text{DM}=56.71 \text{ pc cm}^{-3}$. The supergiant pulse discussed in Cordes et al. (2004) occurs around $t = 50 \text{ s}$. Bottom Frame: The SNR_t vs. m_I is plotted in the top panel. The points associated with the supergiant pulse are plotted in medium gray, points associated with the dispersion measure of the Crab are shown in black, and all other points shown in light gray. A m_I histogram (bottom) plots the total number of events (thick line) and the number of events at $\text{DM}=56.71 \text{ pc cm}^{-3}$ (thin line) as a proxy for the number of pulses detected.

rarer than weaker pulses, Equation 13 tells us that the number of pulses in a m_I bin will decrease as $m_I \rightarrow 0$. There is therefore a tradeoff between choosing a higher $m_{I,\text{max}}$ to allow for the more common, weaker pulses and the false alarm rate.

Figure 7 further explores the role of interstellar scintillations on the modulation index. The SNR_t vs. m_I is plotted for two higher frequencies from Cordes et al. (2004). The data are time-frequency snapshots of pulses determined from previous processing at 1475 and 2850 MHz. The snapshots were only dedispersed at the dispersion measure of the Crab pulsar, the time series were thresholded with $\text{SNR}_{\text{min}} = 3$, and the resulting can-




Figure 7. SNR_t vs m_I of Crab giant pulses at 1475 MHz (top) and 2850 MHz (bottom) from Cordes et al. (2004). The lines represent $m_{I,\text{bb}}$ (Equation 13). The modulation index that divides the pulses from thresholded noise is $m_I \approx 2$ for 1475 MHz and $m_I \approx 1.5$ for 2850 MHz. The modulation indices of the thresholded noise are consistent with those of broadband pulses. The modulation indices of the pulses are systematically larger than what is predicted by Equation 13, suggesting the spectra have inherent structure caused by diffractive interstellar scintillations.

didates were run through the cluster algorithm to find the sample with the maximum SNR. Also plotted is the $m_{I,\text{bb}}$ curve given by Equation 13 with $N_\nu = 118$ at 1475 MHz (top) and $N_\nu = 58$ at 2850 MHz (bottom).

In both panels there are two clusters of points: one due to thresholded noise and one due to the pulses. The thresholded noise lies below $m_I = 2$ and $m_I = 1.5$ at 1475 and 2850 MHz respectively. The typical values for the noise are consistent with $m_{I,T}$ and lie along the broadband curve as predicted in Section 2.3. The modulation indices of the pulses are systematically larger than what is predicted, suggesting the pulses have inherent frequency structure and $m_A \neq 0$. The spectra for the strongest pulse at each frequency are shown in Cordes et al. 2004, Figures 7 and 9, and in both cases they exhibit frequency structure caused by DISS.

4. EXTENSIONS OF THE METHOD

Our discussion has focused the detection of radio bursts using incoherent dedispersion, but the technique has more general applicability. It can also be applied to other classes of signals (periodic objects like pulsars, spectral lines) and other data formats (coherent dedispersion, image cubes from interferometers).

4.1. Periodic Signals

Using the modulation index to characterize the frequency structure of a single pulse can be extended to periodic signals. The procedure is best illustrated by considering a dedispersed time series $\bar{I}(t, DM)$ that has been folded at some period P . A pulse in this folded time series is the sum of N_p pulses

$$\bar{I}_p(t_o, DM) = \frac{1}{N_p} \sum_{t_p} \bar{I}(t, DM) \quad (21)$$

where $t_p = t_o + nP$, t_o is the time sample of the candidate pulse in the folded time series, and $n = 0 \dots N_p - 1$. The modulation index is then calculated for samples where $\bar{I}_p(t, DM)$ is above the SNR threshold. As discussed in Section 2.1 for smoothing in time, calculating the second moment requires that spectra are first summed in time before averaging over frequency, making it obligatory that one returns to the time-frequency data

$$\bar{I}_p^2(t_o, DM) = \frac{1}{N_\nu} \sum_\nu \left[\frac{1}{N_p} \sum_{t_p} I(t, \nu, DM) \right]^2. \quad (22)$$

One can also calculate the second moment for the Fourier analysis method of finding pulses signals but that is beyond of the scope of this paper.

The additional computational cost required to calculate the modulation index can be estimated in the same manner as for the single pulse detections described in Section 3.2. But as a pulsar is a repetitive source, the time span of the data that would need to be reprocessed in a second pass would be larger than for a single pulse. In the extreme case where $N_s = N_t$, the processing load is comparable when the number of candidates is of the order of the number of trial DMs. It is therefore prudent to reduce the number of candidates by other means first before calculating the modulation index of candidate pulsar.

4.2. Coherent Dedispersion

Coherent dedispersion operates on the basedband voltage directly rather than on intensity-like quantities at the output of a spectrometer, which our previous discussion has centered on. In a survey, a set of trial values for DM would be used, as with post-detection dedispersion described earlier. With the correct value of DM, coherent dedispersion restores the phase of the signal to what it was prior to alteration by the ISM (so long as multipath scattering is not important). The procedure deconvolves a phase function from voltage data $\epsilon(t)$ by applying a complex phase filter (e.g., Hankins and Rickett 1975) to produce the dedispersed voltage $\epsilon_d(t)$. The detected signal $I_d(t) = |\epsilon_d(t)|^2$ would then be analyzed with a SNR threshold to identify candidate bursts.

If the true signal is an unresolved pulse with width $W \ll B^{-1}$ where, as above, B is the total bandwidth, the dedispersed output will have a width $W_t \approx B^{-1}$ only if the signal is unmodulated across the band. This requires that any scintillation modulation have a characteristic bandwidth $\gg B$. In this instance, no new information is gained by analyzing the data in the combined frequency-time plane.

For signals that have broader time extent, however, the analysis presented in the paper for post-detection dedispersion still applies. The dynamic spectrum would be calculated for each burst identified in $I_d(t)$ from short time blocks of length Δt_s using the discrete Fourier transform (DFT) $\hat{\epsilon}(t)$ of the baseband signal. Using an N_ν -point DFT, the uncertainty relation implies $\Delta t_s B / N_\nu = 1$. For burst widths $W \gtrsim \Delta t_s$, the series of spectra that span the burst can be used to calculate the modulation index across frequency, which is then used to classify signals as before.

4.3. Application to Images Obtained On a Frequency-Time Grid

The modulation index as discussed so far is calculated for a set of intensity measurements sampled in time and frequency at a particular value of the dispersion measure, $I(t, \nu, DM)$. For signals that are not inherently narrow in time (i.e., not pulsed or transient in nature), no dedispersion is required. Instead, we simply calculate the modulation index m_I as the normalized variance of the intensity across frequency at a particular time,

$$m_I^2 = \frac{\bar{I}^2(t) - \bar{I}(t)^2}{\bar{I}(t)^2}. \quad (23)$$

The modulation index m_I and fractional correlation bandwidth FCB calculated in this manner can be used, for example, to characterize radio observations of a spectral line or maser source (which show coherent frequency structure) and discriminate them from radio frequency interference with a more random frequency structure. We discuss the application of a time-domain modulation index to these source types in Section 5.

The method can be further generalized to apply to interferometric imaging observations of continuum sources that have been acquired in multi-channel modes. While past imaging observations with the Very Large Array (VLA), for example, have only employed a small number of channels (e.g., $\lesssim 32$) or even just one channel in continuum mode, observations with the Expanded Very Large Array (EVLA) or the future SKA pathfinders such as Australian Square Kilometre Array (ASKAP; Johnston et al. 2009) and MeerKAT (Booth et al. 2009) will typically employ many more channels. At the EVLA, the wide bandwidth available for continuum imaging (e.g., 1–2 GHz or 2–4 GHz) requires the use of many spectral channels to avoid chromatic aberration (“bandwidth smearing”). If we were to require a maximum tolerable peak response reduction of 20% at A-array, 1–2 GHz, that would necessitate a channel bandwidth $\Delta\nu$ such that

$$\frac{\Delta\nu}{1.5 \text{ GHz}} \frac{\theta}{\theta_0} < 1,$$

where θ/θ_0 is the source offset from the phase tracking center in units of the synthesized beam.² To maintain sensitivity at that level over just half the primary beam field of view, $\theta/\theta_0 \sim 0.5 \times d/D$ where d and D are the dish diameter and array size respectively. For the EVLA in A-array where $d = 25$ m and $D \approx 36$ km, this constraint requires as many as 720 channels for continuum imaging.

One of the challenges faced by automated source extraction pipelines working on these data is to discriminate between low significance detections of compact sources and random intersections of the point spread function sidelobes caused by strong sources in the field of view. While the exact shape of the sidelobes depends on the details of the array and the observation, they will scale as $I(\theta, \nu) \propto \nu^{-1}$ at an angular distance θ from a bright source. Therefore, an intersection of sidelobes from different sources will show smooth structure in amplitude as a function of channel frequency, while

² See, e.g., the EVLA Observational Status Summary, <http://evlaguides.nrao.edu/>

a true compact source will not. Calculating the modulation index and FCB will allow us to exploit the extra information in the frequency channelization of synthesis images and add an extra discriminant that is easily implementable in automatic pipelines.

4.4. Realtime RFI excision

The modulation index could be useful for systems performing RFI excision on the fly. Interest in realtime incoherent dedispersion transient searches is growing thanks to the large speed-ups achieved by graphics processing units (GPUs) (e.g. Magro et al. 2011). As more of the signal processing moves into hardware, realtime RFI rejection is also moving into instruments. For example Deller (2010); Wayth et al. (2011) have incorporated the calculation of kurtosis into the F-stage of the DiFX software correlator, and the modulation index could be calculated in an identical manner.

In the time-frequency domain the modulation index would be most useful for identifying broadband, impulsive RFI. A sample that had an anomalously low modulation index could be flagged or blanked by the hardware. The modulation index alone should not be used to identify narrowband RFI, as a dispersed pulse is narrowband before the dispersion is correctly accounted for.

5. DISCUSSION

In this paper we have focused our discussion on searching for radio bursts from or similar to those from pulsars and used the modulation index, along with SNR, as a tool for RFI excision. More generally the modulation index can be one among several statistics used to *characterize and classify* signals rather than to filter candidates. Furthermore there may be as-yet undiscovered source classes having different time-frequency signatures.

Low mass stars (M-dwarfs and brown dwarfs) are frequent emitters of radio bursts (Jackson et al. 1989; Berger 2002). Bastian et al. (1990) have shown that the dynamic spectra of dMe flare stars exhibit a rich structure in both time and frequency. The spectra of the flares they observed had both broadband and narrowband components, as well as large spectral indices ($\alpha \approx 10$), suggesting various plasma phenomena at work. A large-scale survey of flare stars could use the modulation index as one of several parameters used to automatically classify flare types. Furthermore, different plasma processes result in different levels of polarization, so calculating the modulation index of each polarization separately could further aid in classifying flares and recognizing RFI.

The study of radio variability in the solar system is a natural extension to flare stars and is another area where the modulation index could be a tool for classifying different radio bursts. Planetary auroral radio emission (ARE) has been detected for all of the magnetic planets in our solar system (Zarka 1998), and dynamic spectra of ARE show complex substructure in both time and frequency. A notable class of fast radio pulses is Jupiter’s “S-bursts”. These bursts have durations of milliseconds and are the brightest of Jupiter’s ARE emission (Zarka 1998). Radio discharges from lightning on the gas giant planets have durations on the order of 10 to 100 ms and fluxes easily detectable at Earth (Zarka et al. 2004). In fact Ruf et al. (2009) have used multi-moment techniques (total intensity and kurtosis) to look for lightning

on Mars during a dust storm. By extension it is expected that extrasolar planets would have radio emission that would be variable and include bright bursts (Farrell et al. 1999; Griebmeier et al. 2007). A study of the modulation indices of flares from Jupiter could help identify similar flares from exo-Jupiters.

Extragalactic sources of fast radio transients might include merging neutron star–neutron star or neutron star–black hole binaries (Hansen and Lyutikov 2001). Short gamma-ray bursts (SGRB) are thought to arise from the merger of two neutron stars, and Pshirkov and Postnov (2010) suggest that GRB monitors could alert low frequency radio observatories (e.g., LOFAR) to a possible radio transient.

Our focus has been on signals that are broad in frequency but narrow in time. The reciprocal problem of spectral lines that are steady in time can be handled in the same way but with the roles of frequency and time reversed. As an example, surveys for maser lines will most likely require a line amplitude that is relatively steady over time scales of days or less. Interstellar scintillation may induce time variations in some maser sources if they are compact enough. SETI (search for extraterrestrial intelligence) often postulates narrowband beacon signals that are constant in time. Interstellar scintillation will certainly induce time variations owing to the compact nature of any relevant transmitters (Cordes and Lazio 1991) but with a modulation index ~ 1 . The methods outlined here apply directly to these problems.

6. CONCLUSION

We have discussed how defining detection statistics based on higher order moments can improve the success of source detection pipelines and focused on the spectral modulation index. By calculating the second moment, we are able to classify signals based not only on their strength but also on their fractional frequency variation. By applying prior information about our target sources, i.e., that they are broadband, the modulation index easily distinguishes between broadband and narrowband signals and allows us to filter a large fraction false positives due to narrowband RFI.

These detection statistics (e.g., SNR, m_1) are crucial to source extraction pipelines because they can be calculated automatically. As new observatories generate more and more data, it is critical that source extraction occurs reliably with minimal human intervention. Although we limit ourselves to two statistics in this paper, a pipeline could make use of many higher-order statistics (e.g. kurtosis, Nita et al. 2007) along with a weighted voting scheme to classify signals in more sophisticated and nuanced ways. For example the modulation index could become one parameter used by detection pipeline based on an artificial neural network (Eatough et al. 2010).

We thank the reviewer for his or her useful comments. This work was supported by NSF grant AST - 1008213. L.G.S. and J.S. were supported by the NASA/New York Space Grant Consortium under grant NNX10AI94H. L.G.S. also received support from the National Astronomy and Ionosphere Center (NAIC). The Arecibo Observatory is operated by SRI International under a cooperative agreement with the National Science Foundation (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana, and the Universities Space

Research Association.

REFERENCES

- M. A. McLaughlin, A. G. Lyne, D. R. Lorimer, M. Kramer, A. J. Faulkner, R. N. Manchester, J. M. Cordes, F. Camilo, A. Possenti, I. H. Stairs, et al., *Nature* **439**, 817 (2006), arXiv:astro-ph/0511587.
- J. M. Cordes and M. A. McLaughlin, *ApJ* **596**, 1142 (2003), arXiv:astro-ph/0304364.
- M. A. McLaughlin and J. M. Cordes, *ApJ* **596**, 982 (2003), arXiv:astro-ph/0304365.
- J. M. Weisberg, B. K. Armstrong, P. R. Backus, J. M. Cordes, V. Boriakoff, and D. C. Ferguson, *AJ* **92**, 621 (1986).
- L. L. Kedziora-Chudczer, D. L. Jauncey, M. H. Wieringa, A. K. Tzioumis, and J. E. Reynolds, *MNRAS* **325**, 1411 (2001), arXiv:astro-ph/0103506.
- S. R. Spangler and L. G. Spitler, *Physics of Plasmas* **11**, 1969 (2004).
- D. R. Lorimer, J. A. Yates, A. G. Lyne, and D. M. Gould, *MNRAS* **273**, 411 (1995).
- M. de Vos, A. W. Gunst, and R. Nijboer, *IEEE Proceedings* **97**, 1431 (2009).
- C. J. Lonsdale, R. J. Cappallo, M. F. Morales, F. H. Briggs, L. Benkevitch, J. D. Bowman, J. D. Bunton, S. Burns, B. E. Corey, L. Desouza, et al., *IEEE Proceedings* **97**, 1497 (2009), 0903.1828.
- J. M. Cordes and B. J. Rickett, *ApJ* **507**, 846 (1998).
- J. M. Cordes, N. D. R. Bhat, T. H. Hankins, M. A. McLaughlin, and J. Kern, *ApJ* **612**, 375 (2004), arXiv:astro-ph/0304495.
- B. J. Rickett, T. H. Hankins, and J. M. Cordes, *ApJ* **201**, 425 (1975).
- J. M. Cordes, *ApJ* **210**, 780 (1976).
- D. R. Thompson, K. L. Wagstaff, W. F. Brisken, A. T. Deller, W. A. Majid, S. J. Tingay, and R. B. Wayth, *ApJ* **735**, 98 (2011), 1104.4900.
- J. S. Deneva, J. M. Cordes, M. A. McLaughlin, D. J. Nice, D. R. Lorimer, F. Crawford, N. D. R. Bhat, F. Camilo, D. J. Champion, P. C. C. Freire, et al., *ApJ* **703**, 2259 (2009), 0811.2532.
- T. H. Hankins and B. J. Rickett, in *Methods in Computational Physics. Volume 14 - Radio astronomy*, edited by B. Alder, S. Fernbach, & M. Rotenberg (1975), vol. 14, pp. 55–129.
- S. Johnston, I. J. Feain, and N. Gupta, in *The Low-Frequency Radio Universe*, edited by D. J. Saikia, D. A. Green, Y. Gupta, & T. Venturi (2009), vol. 407 of *Astronomical Society of the Pacific Conference Series*, pp. 446–+, 0903.4011.
- R. S. Booth, W. J. G. de Blok, J. L. Jonas, and B. Fanaroff, *ArXiv e-prints* (2009), 0910.2935.
- A. Magro, A. Karastergiou, S. Salvini, B. Mort, F. Dulwich, and K. Zarb Adami, *ArXiv e-prints* (2011), 1107.2516.
- A. Deller, in *Proceedings of the RFI Mitigation Workshop. 29-31 March 2010. Groningen, the Netherlands* (2010), p. PoS(RFI2010)035, <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=107>, 1012.0325.
- R. B. Wayth, W. F. Brisken, A. T. Deller, W. A. Majid, D. R. Thompson, S. J. Tingay, and K. L. Wagstaff, *ApJ* **735**, 97 (2011), 1104.4908.
- P. D. Jackson, M. R. Kundu, and S. M. White, *A&A* **210**, 284 (1989).
- E. Berger, *ApJ* **572**, 503 (2002), arXiv:astro-ph/0111317.
- T. S. Bastian, J. Bookbinder, G. A. Dulk, and M. Davis, *ApJ* **353**, 265 (1990).
- P. Zarka, *J. Geophys. Res.* **103**, 20159 (1998).
- P. Zarka, W. M. Farrell, M. L. Kaiser, E. Blanc, and W. S. Kurth, *Planet. Space Sci.* **52**, 1435 (2004).
- C. Ruf, N. O. Renno, J. F. Kok, E. Bandelier, M. J. Sander, S. Gross, L. Skjerve, and B. Cantor, *Geophys. Res. Lett.* **361**, L13202 (2009).
- W. M. Farrell, M. D. Desch, and P. Zarka, *J. Geophys. Res.* **104**, 14025 (1999).
- J.-M. Grießmeier, P. Zarka, and H. Spreew, *A&A* **475**, 359 (2007), 0806.0327.
- B. M. S. Hansen and M. Lyutikov, *MNRAS* **322**, 695 (2001), arXiv:astro-ph/0003218.
- M. S. Pshirkov and K. A. Postnov, *Ap&SS* **330**, 13 (2010), 1004.5115.
- J. M. Cordes and T. J. Lazio, *ApJ* **376**, 123 (1991).
- G. M. Nita, D. E. Gary, Z. Liu, G. J. Hurford, and S. M. White, *PASP* **119**, 805 (2007).
- R. P. Eatough, N. Molkenhuth, M. Kramer, A. Noutsos, M. J. Keith, B. W. Stappers, and A. G. Lyne, *MNRAS* **407**, 2443 (2010), 1005.5068.